

<b>Report Documentation Page</b>			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE <b>01 OCT 2014</b>	2. REPORT TYPE <b>N/A</b>	3. DATES COVERED <b>-</b>		
4. TITLE AND SUBTITLE <b>Contract-Based Integration of Cyber-Physical Analyses</b>			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) <b>Ivan Ruchkin, Dionisio De Niz, Sagar Chaki, David Garlan</b>			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213</b>			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>				
13. SUPPLEMENTARY NOTES <b>The original document contains color images.</b>				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>SAR</b>	
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>	18. NUMBER OF PAGES <b>1</b>	19a. NAME OF RESPONSIBLE PERSON

# Contract-Based Integration of Cyber-Physical Analyses

Ivan Ruchkin, Dionisio De Niz, Sagar Chaki, David Garlan. Carnegie Mellon University, Pittsburgh, PA, USA.

## Problem

CPS engineering combines diverse model-based *analyses* from various engineering domains. Differences in domain abstractions lead to integration issues:

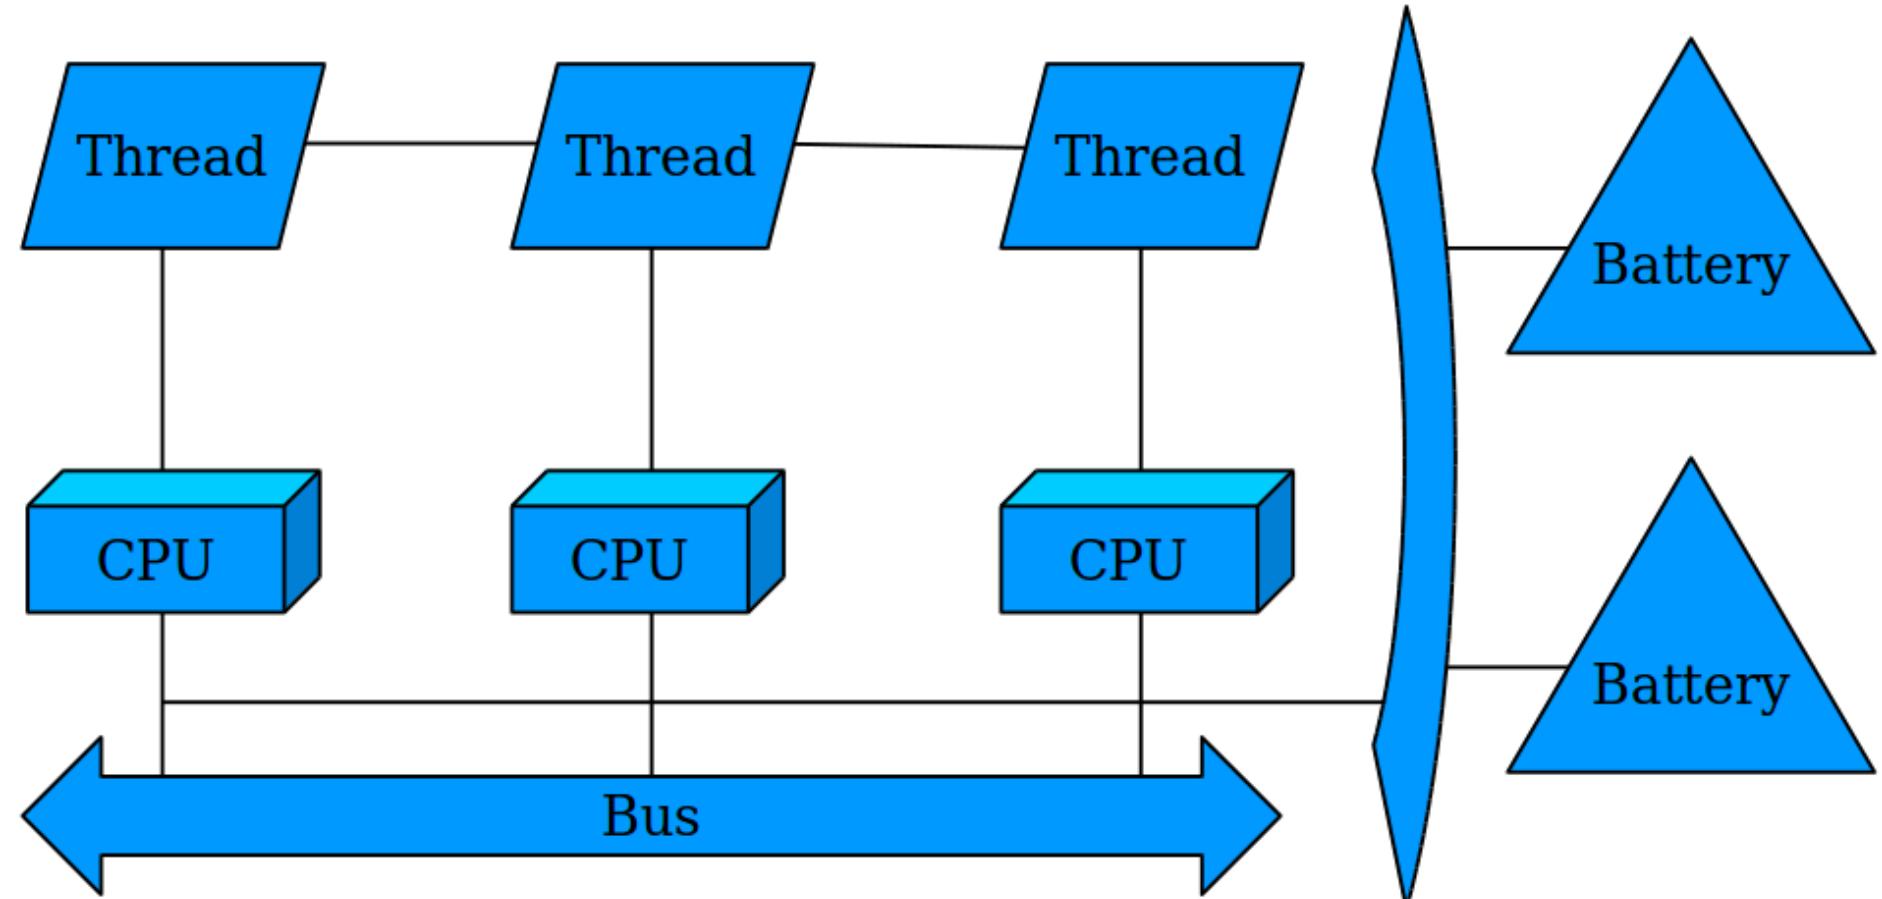
- If an assumption of an analysis are violated by another, the outputs of the former may be invalid.
- Specification of such implicit assumptions and detection of their violation is left to human designers, who are often unable to cope with complexity.
- Analysis integration problems discovered late in development lead to expensive changes to the system.

Hence the research question:

- How to specify analysis compositions and verify their correctness?

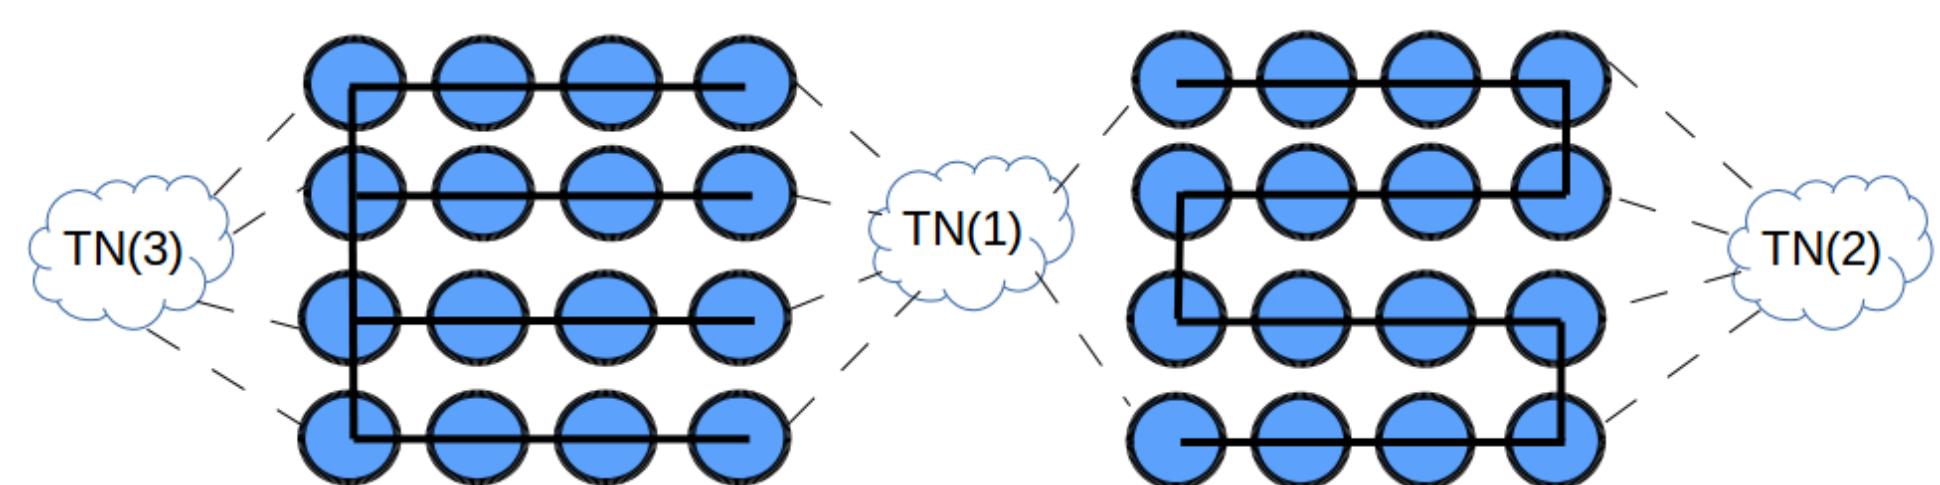
## Example System

Consider an autonomous aircraft as an example CPS. It operates data with different classes of security, from normal to top secret (ThSecCl). Periodic threads ( $T$ ) execute on several processors ( $C$ ). The aircraft is powered with multi-cell reconfigurable batteries ( $B$ ). The system's architecture shown below is specified in AADL.



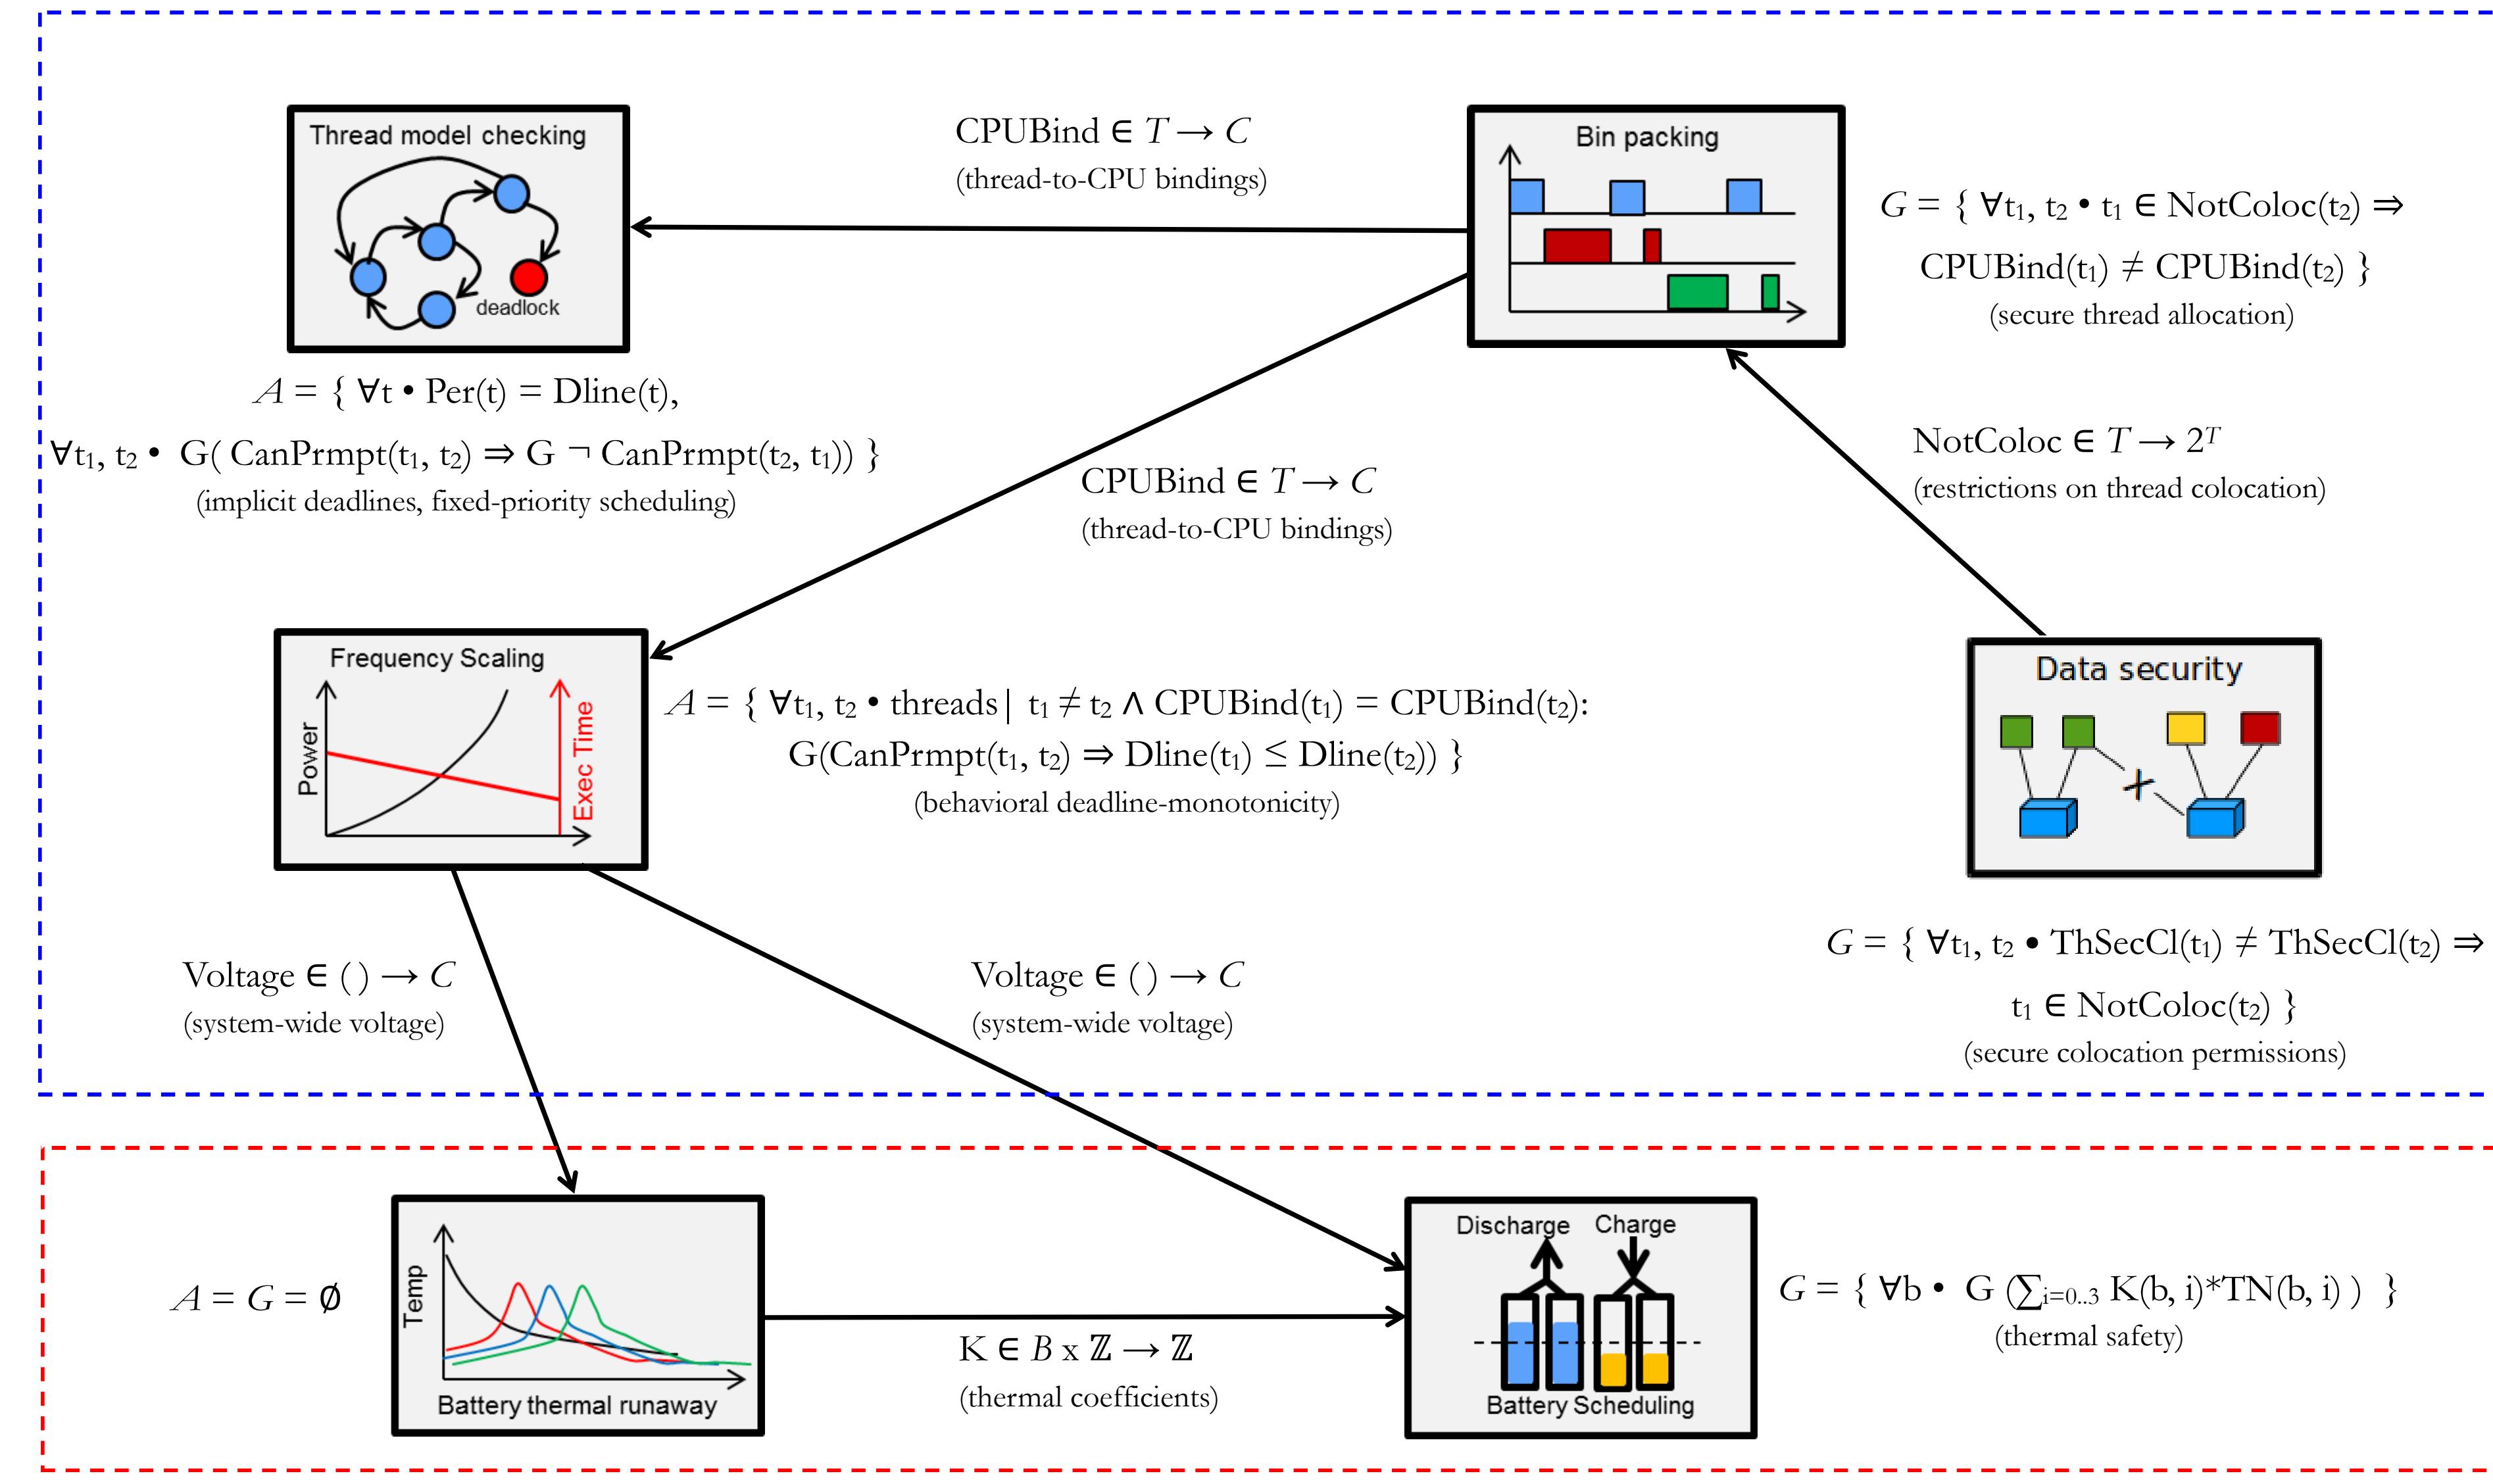
A battery has a matrix of cells, and each cell has a current level of charge. A battery scheduler determines parallel and sequential connections between groups of cells in order to satisfy voltage and current output requirements.

Thermally, each cell exchanges heat with its neighboring cells (*thermal neighbors*, TN) through an electrical connector, affecting the risk of a *thermal runaway*.



## Example Analyses

Scheduling verification domain  $\sigma_{\text{sched}}$



## Verification Domains

A verification domain  $\sigma = (\mathcal{A}, \mathcal{S}, \mathcal{R}, \mathcal{T}, \llbracket \cdot \rrbracket_\sigma)$  formalizes domain-specific constructs for several related analyses.

- $\mathcal{A}$  — a set of sorts, comprised of system elements and standard sorts. E.g., integers  $\mathbb{Z}$ , threads  $T$ , or scheduling policies  $\text{SchedPol}$ .
- $\mathcal{S}$  — a set of static functions that encode design-time properties. E.g., thread period  $\text{Per}$ , thread-to-CPU binding  $\text{CPUBind}$ , and system-wide  $\text{Voltage}$ .
- $\mathcal{R}$  — a set of runtime functions that encode dynamic properties. E.g., preemption relation  $\text{canPrmp}(t_1, t_2)$  and number of cells in a battery  $b$  with  $i$  thermal neighbors  $\text{TN}(b, i)$ .
- $\mathcal{T}$  — execution semantics of  $\sigma$  — a set of sequences of assignments to  $\mathcal{R}$ . We use Promela programs to implement the semantics.
- $\llbracket \cdot \rrbracket_\sigma$  — a domain interpretation of  $\mathcal{A}$ ,  $\mathcal{S}$ , and  $\mathcal{T}$ . E.g.,  $\llbracket \text{SchedPol} \rrbracket_\sigma = \{\text{RMS}, \text{DMS}, \text{EDF}\}$ .

Formally, an AADL architectural model  $\mathbf{m}$  is an interpretation  $\llbracket \cdot \rrbracket_\mathbf{m}$  of  $\mathcal{A}$ ,  $\mathcal{S}$ , and  $\mathcal{T}$ . E.g.,  $\llbracket T \rrbracket_\mathbf{m} = \{\text{SensorSample}, \text{Ctrl}_1, \text{Ctrl}_2\}$ ,  $\llbracket \text{CPUBind} \rrbracket_\mathbf{m} = \{(\text{Ctrl}_1, \text{CPU}_1), (\text{Ctrl}_2, \text{CPU}_2), \dots\}$ .

$\llbracket \cdot \rrbracket_\sigma \cup \llbracket \cdot \rrbracket_\mathbf{m}$  form a full interpretation of  $\mathcal{A}$ ,  $\mathcal{S}$ ,  $\mathcal{R}$ , and  $\mathcal{T}$ .

## Analysis Contracts

Each analysis is assigned a *contract* — a tuple  $(I, O, A, G)$ .

- Inputs  $I \subseteq \mathcal{A} \cup \mathcal{S}$  declare elements that the analysis reads.
- Outputs  $O \subseteq \mathcal{A} \cup \mathcal{S}$  declare elements that the analysis writes.
- Assumptions  $A \subseteq \mathcal{F}_\sigma$  are logical statements that must be satisfied by every input model to the analysis:  $\mathbf{m} \models \varphi$ .
- Guarantees  $G \subseteq \mathcal{F}_\sigma$  are logical statements that must be satisfied by every output model of the analysis:  $\mathbf{m} \models G$ .

Assumption and guarantee formulas have the following syntax:

$$\mathcal{F}_\sigma ::= \forall v_1 \dots v_j \cdot \varphi \mid \exists v_1 \dots v_j \cdot \varphi \mid \\ \forall v_1 \dots v_j \cdot \varphi : \psi \mid \exists v_1 \dots v_j \cdot \varphi : \psi,$$

where  $\varphi$  is a predicate logic formula over  $\mathcal{A} \cup \mathcal{S}$ ,  $\psi$  is an LTL formula over  $\mathcal{A} \cup \mathcal{S} \cup \mathcal{R}$ .

## Analysis Ordering

Correct execution of analyses requires satisfaction of all input-output dependencies for each analysis. Formally, contract  $C_i$  depends on contract  $C_j$  if  $C_i.I \cap C_j.O \neq \emptyset$ .

An ordering  $\langle C_1 \dots C_n \rangle$  of contracts is sound if and only if predecessors are not dependent on successors:

$$\forall i \in [1, n] \cdot \forall j \in [1, i] \cdot C_j.I \cap C_i.O = \emptyset.$$

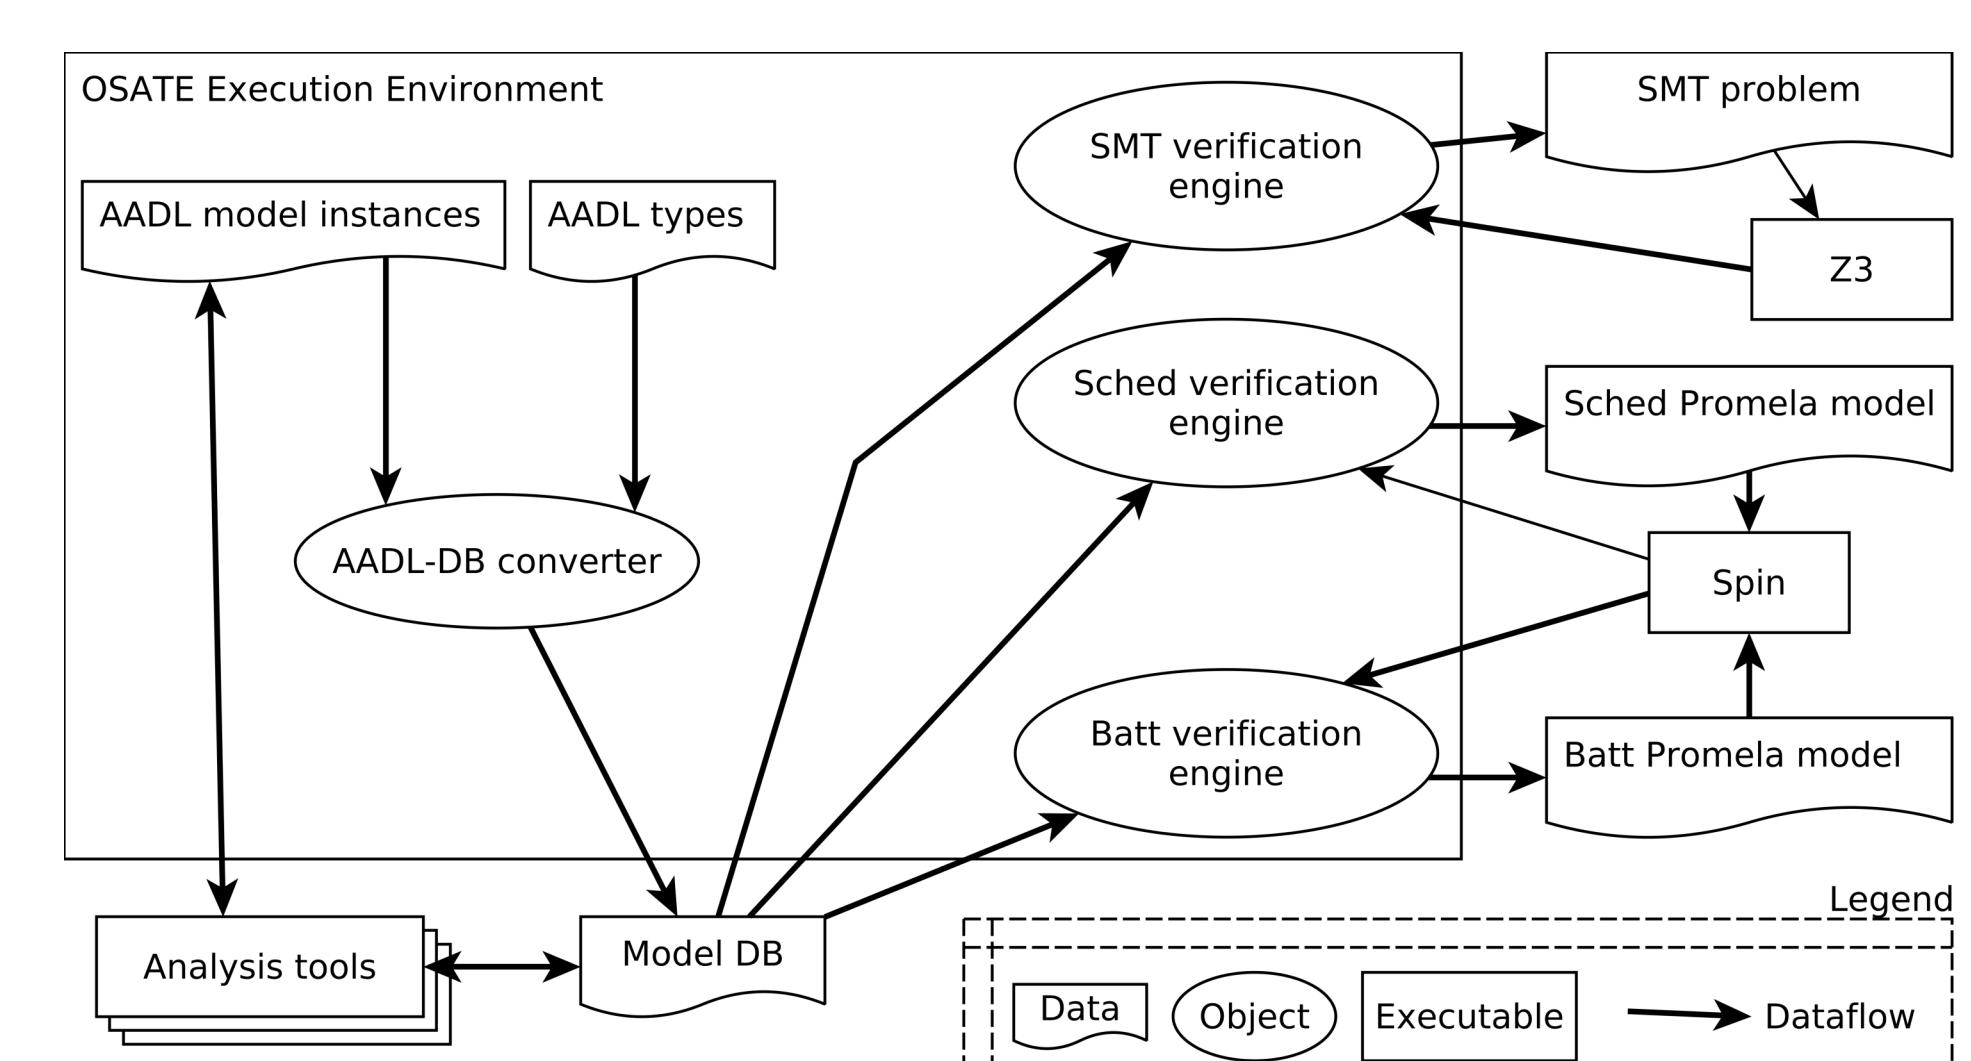
Consider a graph with vertices being contracts and edges being contract dependencies. There exists a sound ordering of contracts if and only if the graph is not cyclic. If it is not cyclic, any topological ordering is sound.

## Contract Verification

The goal of contract verification is to decide  $\mathbf{m} \models \mathcal{F}_\sigma$ .

For purely first-order formulas that contain only  $\varphi$ , we decide satisfiability via SMT solving. An SMT program is generated based on  $\mathcal{A}$  and  $\mathcal{S}$  mentioned in  $\varphi$ , and an SMT solver is invoked on  $\neg \varphi$  (or  $\varphi$  for existential quantification). A universally (existentially) quantified contract is satisfied if and only if UNSAT (SAT) is returned.

For formulas combining predicate formula  $\varphi$  and LTL formula  $\psi$ , we first generate an SMT program for  $\varphi$  and find all valuations of  $v_1 \dots v_j$  that satisfy  $\varphi$ . For each such valuation we call Spin on a Promela program that implements  $\mathcal{T}$  for  $\mathbf{m}$  in the domain of  $\psi$ . Formula  $\psi$  is transformed into an LTL property specification in Promela. A universally (existentially) quantified contract is satisfied if and only if the LTL property holds for all (at least one) valuations. The architecture of our verification tool is shown below:



## Experimental Results

**Effectiveness:** we have been able to detect analysis integration errors and verify their absence for each analysis in the example.

**Scalability:** the results of scalability experiments with our implementations of  $\mathcal{T}$  are shown in the tables below.

$\mathcal{T}_{\text{sched}}$ :

Threads	(R/D)MS time*	EDF time*	Cells	FGURR time*	FGWRR time*	GPWRR time*
3	0.01	0.01	9	0.13	0.15	0.15
4	0.01	0.52	12	0.61	2.34	3.94
5	0.07	33.4	16	44	31.4	127
6	0.37	2290.0	20	1060	619	memlim
7	2.18	memlim	8	12.4	memlim	memlim
9	71.2	memlim	9	71.2	memlim	memlim
10	421	memlim	10	421	memlim	memlim
11	memlim	memlim	11	memlim	memlim	memlim

\* All times are in seconds.

Copyright 2014 ACM. This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. NO WARRANTY: THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT. This material has been approved for public release and unlimited distribution. Carnegie Mellon® is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University. DM-0001715.